arXiv:1405.3496v1 [cs.SE] 14 May 2014

Current Concepts in Version Control Systems

Petr Baudiš*

2009-09-11

Abstract

We give the reader a comprehensive overview of the state of the Version Control software engineering field, describing and analysing the concepts, architectural approaches and methods researched and included in the currently widely used version control systems and propose some possible future research directions.

Preface

This paper has been originally written in May 2008 as an adaptation of my Bachelor thesis and I have been meaning to get back to it ever since, to add in some missing information, especially on some of my research ideas, describing the patch algebra more formally and detailing on the more interesting chapters of historical development in the field.

Unfortunately, it is not likely that I will ever get the time to polish the paper up to my full content; in the meantime, it is slowly getting obsolete and I believe it can have a good value even as it stands, especially as an introductory material for people getting into this area. Maybe I have missed some of the latest developments; I ask the reader for tolerance of the omissions, and my fellow researches and authors for followups on this work that would fill in the gaps and explain ongoing new ideas in the field.

— Petr Baudiš

1 Introduction

Version control plays an important role in most of the creative processes involving a computer. Ranging from simple undo/redo capabilities of most office tools to complicated version control supporting branching and diverse collaboration of thousands of developers on huge projects. This work will focus on the latter end of the scale, describing the modern approaches for advanced version control with emphasis on software development. We will compare various systems and techniques, but not judge their general suitability as the set of requirements for version control can vary widely. We will not focus on integration of version control in the source configuration management systems[1] since SCM is an extremely wide area and many other papers already cover SCM in general. [3] [15] [70]

This work aims to fill the sore need for an introduction-level but comprehensive summary of the current state of art in version control theory and practice. We hope that it will help new developers entering the field to quickly understand all the concepts and ideas in use. Since the author is (was) a Git developer, we focus somewhat on the Git version control system in our examples, but we try to be comprehensive and cover other systems fairly as much as is within our power and knowledge.

We also describe for the first time some techniques so far present only as source code in various systems without any explicit documentation and present some of our original work. One problem in the field of version control is that since most innovation and research happens not in the academia or even classical industry, but the open community, very few of the results get published in the form of a scientific paper or technical report and many of the concepts are known only from mailing list posts, random writeups scattered around the web or merely source code of the free software tools. Thus, another aim of this work is tie up all these sources and provide an anchor point for future researchers.

*The work on this paper was in part sponsored by Novell (SUSE Labs).

[1] Source Configuration Management covers areas like version control, build management, bug tracking or development process management.

2 Basic Design

The basic task of version control system is to preserve history of a collection of files (*project*) so that the user can compare two historical versions of the content easily, return to an older version or review the list of changes between two versions.

2.1 Development Process Model

We will focus on version control systems that are based on the *distributed development* paradigm. While the concept dates many years back, distributed development has only recently seen major adoption and so far mostly only in the area of open source development.

The classical approach to version control is to have a single central server[2] storing the project history; the client can have the current version cached locally, and possibly also (parts of) the history, but any operation involving change in the history (e.g. committing new version) requires cooperation with the central server.

This *centralized development* approach has two major shortcomings: first, the developer has to have network access during development, which is problematic in the era of mobile computing — i.e., when developing during travel or fixing problems at the customer site. Second, the central repository access implies some inherent bureaucracy overhead — the developer has to gain appropriate commit permissions before being able to version-control any of his work, development of third-party patches is difficult and since the repository is shared, the developer is usually restricted in experimentation and creating short-lived branches, often resorting to using some kind of manual version control during regular workflow and using the VCS only on official occassions.

On the other hand, since all the development is tracked at a central place, this model may pose advantages to corporate shops because work and performance of all developers can be tracked easily.[3]

In the *distributed development* model, a copy of the project created by version control has full version tracking capabilities and not only is all the history usually copied over (*pulled*), but it is also possible to commit new changes locally; later, they can be sent (*pushed*) to another repository in order to be shared with other developers and users. This approach effectively makes every copy a stand-alone repository with stand-alone branches; nevertheless, most systems make it easy to emulate the centralized version control model as a special-case at least for the read-only users.

2.2 Snapshot-Oriented History

There are two possible approaches to modeling the history. The first-class object is usually the particular *version*[4] as a state (*snapshot*) of the content at a given moment, while the changes are deduced only by comparing two versions. However there is also a competing approach which focuses on the *changes* instead: in that case, the important thing when marking new version is not the new state of the content, but the difference between the old and new state; a particular version is then regarded as a certain combination of changes.[5]

2.2.1 Branches

With regard to the snapshot-oriented history model, it may suffice in the trivial case to merely represent the succeeding versions as a list with the versions ordered chronologically. But over time it may become feasible to work on several versions of the content in parallel, while also recording further modifications using the version control system; the development has effectively been split to several variants (*branches*).

Handling branches is an obvious feature of a version control system as the necessity to branch the development is frequent for many projects targetted at wider audience. Commonly, aside of the main development branch another branch is kept based on the latest public release of the project and it is dedicated to bugfixes for this release.

Moreover, branches have many other uses as well — it may be feasible for individual developers or teams to store work in progress yet unsuit-

[2] Possibly with several read-only mirrors for replication, as in the BSD ports setup.

[3] In case of in-company development using distributed version control, this can be alleviated by setting appropriate policy for the developers to push their work to designated central repositories in regular intervals.

[4] We use the terms *version*, *revision* and *commit* interchangeably in this paper based on the context, according to usual practice.

[5] Note that the conceptual model may have no relation to actual storage, which can well use deltas even within snapshot-oriented systems.

able for the main development branch in separate branches, or even for a developer to save incomplete modifications during short-time reassignment to a different task or when moving to another machine. The prerequisite for this is that creating a separate branch must have only minimal overhead, both technical (time and space required to create a new branch) and social (it must be possible to easily remove the branch too; also, the unimportant branches should not cloud the view of the important major ones).

The ability to properly record and manipulate branches nevertheless becomes truly vital as the development gets distributed, since in effect every clone of the project becomes a separate branch. If the system aims to fully and comfortably support the distributed development model, it must make sure that there are no hurdles for creating many new branches; furthermore, there can be no central registry of branches, thus the branch-related data structures must allow for branches to grow from different points in history with full independence and dealing with naming conflicts,[6] shall they ever meet again in some repository.

Aside of the obvious operations with branches (referring to them, comparing them, storing them in a reasonable manner, switching between them), perhaps the technically most challenging one is merging — combining changes made separately in two branches in a single revision. In the past, only the most crude methods for merging two branches were available, but again with the onset of distributed version control this aspect has grown in importance (a merge essentially happens every time a developer exchanges modifications with the outside world). We will explain the challenges associated with merging in section 6; one important consideration we shall mention now is that many of the merging techniques require to find the latest point in history common to both of the to-be-merged branches (or, more generally, which changes are not common to the history of both branches).

Aside of merging, another desirable operation is so-called *cherry-picking*: taking individual changes from one branch to another without merging them completely. While most modern systems support this in a rudimentary way, the operation is not natural in snapshot-oriented systems and can result in trouble when merging the branches later, both in duplicated changes stored in the history and increased merge conflicts rate.[7]

2.2.2 Data Structures

Given that the development may fork at some point in history, we cannot use simple lists anymore to adequately represent this event but we have to model the history as a directed tree instead, with succeeding versions connected by edges; the node representing a *fork point* (not only one, but several branches grow from the node) has multiple children.

However, when we take merging into consideration, it becomes necessary to also represent the past merges between branches if we are to properly support repeated merges and preserve detailed history of both branches across the merge. The commonly used method is to use a more general *directed acyclic graph* (DAG; figure 2 several pages later) to represent the history; compared to a directed tree, DAG lifts the restriction for a node to have only a single parent. Thus, versions representing merges of several branches have the last ("head") versions of all the merged branches as their parents. The acyclic property makes sure that a commit still cannot be its own ancestor, thus keeping the history sane (time travel and quantum effects notwithstanding).

An interesting problem to note is how to actually define a branch. In case of trees, the branch is defined naturally by the tree structure, however in a DAG this becomes less obvious. One approach is to record branch of each revision explicitly (e.g. Mercurial), another is to give up on sorting out revisions to explicit branches altogether (e.g. Git) — in that case, a "branch" is merely a named pointer to the commit at the head of the branch, with no clear distinction of which ancestors belong to the same branch; it might be tempting to keep branch correspondence through fixed parent nodes ordering in merge nodes ("first parent comes from the same branch"), however this approach has caveats as described in 6.1.

[6] In independent repositories, branches and revisions can be given conflicting names, as will be covered in detail later.

[7] Git provides a **git-rerere** tool for automatically resolving merge conflicts that have been resolved in a certain way before, thus somewhat reducing this problem.

2.3 Changes-Oriented History

Alternatively to snapshot-oriented history, a radically different approach is to focus on the changes as the first-class object — the user creates and manipulates not a particular version of the project, but instead a particular *change*; a "version" is then described simply by given combination of changes recorded. The major advantage is that merging operations may be much more flexible, and inherently more data is preserved as the changes themselves are recorded and do not have to be deduced a posteriori from checking the difference between two snapshots.

The main representative of this class is **Darcs** (see 3.7) with its main trait of the powerful formalisation *patch algebra* (see 6.6) which makes it possible to easily perform merging of two versions by adding the changes missing in the other version and reordering the changes in a well defined way.

In order to make full use of the smart merging techniques like patch algebra, in some cases the operation time can degrade exponentially to the number of revisions in the current implementations [21]; making the algorithms more efficient is subject of active research, however at least until very recently, practicality of changes-oriented history model was severely limited and this approach was not suitable for large histories; we are not aware of performance improvement studies for the recent Darcs improvements and it is uncertain if these complexity issues are inherent to the whole changes-oriented model.

2.4 Objects Identification

The most visible (and diverse) aspect of the history models is the problem of identifying the history objects (snapshots or changes, depending on the model).

The old **RCS** and **CVS** tools use a simple scheme where every version is identified by a string of dot-separated numbers; the first number is usually 1, the second number is the version number on the main branch. If it is required to identify a version on a different branch, `BASE_VERSION.BRANCH.VERSION` scheme is used — `BRANCH` is the number of the branch forked from the base version and `VERSION` is the sequence number of the given version on the branch; this scheme can be used recursively to identify versions on branches of branches etc.

However, we shall remember that in a distributed system, we cannot see all the used identifiers in all repositories (since they can be on disconnected systems), thus we meet with special problems when having to simultanously assign new identifiers in different disconnected repositories. Eventually, we are met with three conflicting requirements[8] [44][6]:

- *Uniqueness:* The identifier is guaranteed to be unique and always identifies only a single history object in a given repository. (Note that this requirement is frequently slightly compromised by only hoping that the identifier is unique with a very high probability.) [9]
- *Simplicity:* The identifier is easy to use — reasonably short and preferrably informative to the user (e.g. giving some information on ordering of the revisions).
- *Stability:* The identifier does not change over time and between different repositories.

The requirements are of conflicting nature — we can generally always pick only two and sacrifice the third. Different version control systems make different choices here: for example, Git, Mercurial and Monotone sacrifice some usability and assigns every commit a 40-digit SHA1 hash [57] depending on the current content and history of the version.[10] To alleviate the UI problem somewhat, the systems allow shortcutting the long hashes to only few leading digits, as long as they match uniquely.

[8] For clarity, in this context we use "simple" instead of "memorable" and "stable" instead of "global" presented in the literature. The latter correspondence is somewhat loose and given under the assumption that it does not make sense to use central authority for assigning identifiers in case of distributed version control — thus, either the identifiers are globally-meaningful and stable, or they have only local meaning and thus inter-repository collisions are inevitable and relabeling in that case necessary.

[9] Here we mean uniqueness guaranteed by technical means, preferrably provably secure. E.g., GNU Arch does use identifiers that are "unique", but only by policy as one part of the id is user's e-mail address — these are unique, but there is no technical safeguard that different users will not input the same e-mail address; this approach requires trusting the other repository database to contain correct content.

[10] This also technically sacrifies some uniqueness, but the probability of collision of two hashes is so small that it is deemed to be unimportant for practical purposes. This has been disputed by [5] [34], however that paper has been criticized by the version control [40] and cryptographic [14] community.

BitKeeper and Bazaar[11] sacrifice the stability and use RCS-style identifiers that are unique only within a single repository and can shift during merges. Darcs sacrifies the uniqueness and uses user-assigned identifiers for individual changes (however, the identifiers are encouraged to be rather elaborate, so collision probability should not be very high).

2.5 Content Movement Tracking

One interesting problem in version control is following particular content across file renames, copies, etc. Practically all systems offer a way to list all revisions changing given file or set of files, however things become more troublesome when a file is renamed or copied or when we want to check history of finer-grained objects than whole files. Most systems support so-called "annotated view" of a file, for each line showing the commit that introduced it in its current version. However, this operation is frequently very expensive and on many occassions the line-level granularity is not suitable either.

While old systems like CVS do not support tracking file renames at all and the new files start with fresh history unless the repository is manually modified,[12] most newer systems support file renames, usually by providing a special command that records the movement — less common approach used e.g. by GNU Arch is to autodetect movement by tracking unique tags within the objects (strings within files or hidden files in directories).

In that case, showing history of the new file will usually automatically switch to the older file name at the point of the rename, and there is also a possibility of cross-rename merges when a file has been renamed between the two revisions. However, this carries e.g. the danger of making the merging operation complexity proportional to project history, something many systems seek to avoid. Also, as the support gets naturally extended to recording not only renames but also file copies, it becomes much less obvious which file instances to merge; another natural extension to record file combinations makes the situation even more challenging.[13]

A somewhat controversial alternative approach has been taken by Git, where file renames and copying is not explicitly recorded with the rationale that Git aims to track content on a more fine-grained level and hard-coding renames information would pollute the history with potentially bogus information. [25] Instead, Git can heuristically detect the events by computing similarity indexes between files in the compared revisions,[14] and it provides a so-called "pick-axe" mechanism to show changes adding or removing a particular string. Thus, when feeding pick-axe with some chunk of code, it will show the commits that introduced it to the files, as well as commits that removed it from files, effectively giving the user a rudimental ability to trace the code history.[15]

The problem of showing history of particular parts of files and following semantic units across files is still open and will in our opinion benefit from further research (as elaborated in 7.1).

3 Current Systems

We have already mentioned several version control systems when describing basic design aspects; now, we shall briefly describe the most widespread and interesting ones in more detail. We will focus on systems that are freely available and in use in the open-source community, and even so not go through all of them but instead highlight the most influential ones. We will not cover most proprietary systems like ClearCase or SourceSafe in detail since we feel they generally do not bring much innovation into the version control field itself (focusing more on process management, integration with other software etc.) and detailed information is hard to gather from open sources.

Up to CVS, the evolution was rather slow, but then with growing need for distributed version control, many short-lived projects sprang up

[11] Technically, Bazaar internally uses globally unique identifiers, but RCS-style identifiers are the primary way of user interaction.

[12] The modification usually involves manually copying the file under new name within the repository database. However, this presents problems when the working tree is seeked to an older point in the history or in case of branches.

[13] The underlying design of recording renames varies — for example Subversion records renames as copy+delete operations, which can bring some inherent problems e.g. for using the rename information during merges. [50]

[14] The similarity index is computed by hashing line-based file chunks of the files and comparing how many of the hashes exist in both files. [29]

[15] An obvious idea for an extension is to have the GUI frontends integrate mouse text selection with pickaxe, further simplifying the process.

and explored various approaches; now the field has rather stabilized again, with three major systems currently used: Bazaar, Git and Mercurial.

3.1 The Two Grandfathers

The first true version control ever in use was **SCCS** (Source Code Control System) [63], introduced in 1972 at Bell Labs and distributed commercially. It uses the weave file format (4.2.5) to track individual files seprately and provides only a rather arcane user interface. Although its use in the industry is very rare nowadays, the weave storage format has resurfaced recently.

A competing system **RCS** (Revision Control System) [54] published in 1985 has seen much more success, even though it also operates only on isolated files. It uses a delta sequence file format (4.2.4) and the rapid adoption rate compared to SCCS is likely to be accounted mainly to much more user-friendly interface (though it is still rather rudimentary by current standards). RCS is still sometimes used in present for version-controlling individual files, e.g. when maintaining system configuration. However, its concepts and file format has seen even much more widespread adoption thanks to CVS.

3.2 Concurrent Version System

CVS [17] is essentially an RCS spinoff, bringing two major extra capabilities: ability to version-control multiple files at once and support for parallel work, including network support. It has seen extremely widespread adoption and has long remained the de-facto standard in version control, at least among the free systems.

However, CVS is directly based on RCS, and it is ridden with inherited artifacts of individual file revision control — the need for elaborate locking, difficult handling of branching and merging and lack of commits atomicity: when committed change spans multiple files, the commit will be stored for each file separately and there is no reliable way to reconstruct the whole change.[16]

3.3 Modern Centralized Systems

Perforce [49] is a widely commercially used proprietary version control system that provides rather advanced version control features, though it keeps within the centralized model. Compared to CVS, it supports atomic commits (tying changes in multiple files together), tracking file renames and more comfortable branching and merging support (branches are presented as separate paths in the repository and repeated merges are supported).

In the open-source version control realm, **SVN** [59] has been developed as a direct CVS replacement and currently appears to remain the final step in the evolution of centralized version control. It shares many of the Perforce features, though many workflows are different. It is seeing rapid adoption [58] and basically dethroned CVS as the default choice of centralized version control for open-source projects.

3.4 Older Distributed Systems

BitKeeper [11] is another proprietary version control system worth mentioning, since it pioneered the distributed version control field[17] and also has relatively well-understood capabilities and internal working as it was for long time available for free usage by open source projects and some prominent projects like the Linux Kernel used it.

BitKeeper supports atomic commits (though files have revisions tracked individually as well), generic DAG history and the basic distributed version workflows;[18] however, each repository contains only one branch. RCS-style revision numbers are used, same revisions can have different numbers in different repositories and revision numbers on branches can shift during merges. Internally, BitKeeper uses the weave file format and makes full use of its annotation capabilities in the user interface and during merges (see 6.4). It is still actively developed and reportedly commercially popular.

With regard to the history, we shall also mention **GNU Arch** [32] as the first open-source distributed version control system that has seen wider usage, however it suffered by bad performance and difficult user interface; the project has been basically discontinued (though still formally maintained), but Bazaar (see below) can

[16] However, relatively reliable heuristics are implemented, e.g. by the **cvsps** project. [18]

[17] The concept of distributed version control has been explored before by e.g. Aegis, but BitKeeper was the first system to bring it to widespread use.

[18] That is, the pull/push operations and good merging support.

trace back its ancestry to GNU Arch to a certain degree; Darcs has been also somewhat inspired by its changeset-oriented design.

3.5 Monotone

Another important milestone is **Monotone** [39] with its design built around a graph of objects identified by hash numbers computed from their contents[19] serving as direct inspiration for Git and Mercurial. Monotone itself puts strong emphasis on security and authentication, using RSA certificates during communication, etc. It is actively developed but does not see as wide adoption as the three systems below, mainly due to performance issues.

Git and Mercurial build upon the Monotone object model, which is generally built as follows (Monotone implementation details notwithstanding): There are three main object classes — files (or blobs), manifests (or directory trees) and revisions (or commits/changesets). All objects are referenced by large-size hash of their content that should have guarantees akin to cryptographical hashes (currently, SHA1 is used in all popular systems). File objects contain raw contents of particular file versions. Manifest objects contain listing of all files and directories[20] at a particular point of history, with references to objects with appropriate content. Revision objects form the history DAG as a subgraph of the general object graph by referencing the parent revision (or revisions, in case of merge) and the corresponding manifest object. Usually, log message and authorship information is also part of the revision object.

The reader should notice that this scheme is *cascading* — by changing a single file, new manifest object with unique id will be generated (the file hash inside the manifest object changed, thus manifest object hash will change too), and this will in turn result in a revision object with unique id as well. Even committing the same content with the same log message in different branches will not create an id clash since the parent revision hash will be different within both objects. Tagging can be realized by simply creating named reference to either the revision object or another object containing the reference and also some tag description, possibly digitally signed by the project maintainer to confirm that the tagged revision is an official release; thanks to the usage of cryptographically-strong hashes for references, the cryptographic trust extends to the complete revision content.

[19] Technically, the first system using hashes as identifiers is **OpenCM** [46], but this system has never really taken off as far as we can tell.

[20] Git uses one object per directory instead; there is a trade-off between object reuse and dereferencing overhead in this choice.

3.6 The Three Rulers

In spring of 2005 a rapid series of events resulted in termination of the permission to use BitKeeper for free software projects, providing an immediate impulse for improvement in the area of open-source version control systems, especially to the Linux Kernel community — almost in parallel, the Git and Mercurial projects were started. Independently, Canonical Ltd. (the vendor of popular Linux distribution Ubuntu) started the Bazaar[21] project around the same time.

All three systems provide atomic commits, store history in generic DAG and support basic distributed version control workflows.

Git [24] [25] re-uses the generalized Monotone object model; by default it stores each object in a separate file in the database, thus having a very simple and robust but inefficient data model;[22] however, extremely efficient storage using Git Packs (4.3.2) is also provided. In its user interface, Git emphasizes (but does not make mandatory) the rather unique concept of *index*, providing a staging area above the working tree for composing the next commit.[23] Git is written in C (and a mix of shell and Perl

[21] At that time, it has been called BazaarNG — a successor to an original Baz(aar) project which was itself a fork of GNU Arch.

[22] This scheme precludes need for any locking whatsoever, but there is no compression in use and using many small files means inherent overhead on the filesystem level.

[23] Originally a low-level mechanism, the popularity of index usage among developers elected it to a first-class user interface concept. If users choose to use it, they manually add files (or even just chunks of changes) from the working tree to the index at various points and then the index is committed, regardless of working tree state; thus, the user may modify a file, add it to index and then before committing modify it again for example with the changes required to compile it locally — the commit operation will use the older version of the file. When explicitly recording all files to commit in the index, greater discipline of making appropriately fine-grained commits can be encouraged. Also, non-trivial conflicts are recorded in the index, making for a natural way of resolving them by simply adding the desired resolution to the index.

scripts) and is tightly tailored to POSIX systems, which historically meant portability problems especially with regard to Windows. It is designed as a UNIX tool, providing large number of high- and low-level commands, and it is traditionally extended by tools calling the commands externally.

Mercurial [37] [68] has very similar object model to Git and Monotone, but uses the revlog format (4.3.3) for efficient storage. It is written mainly in Python and thus does not suffer from the portability problems of Git; since it uses an interpreted language, the main extension mechanism is writing tightly integrated plugins. Performance-sensitive portions are implemented in C.

While both of the systems above always emphasized their robust data models and performance, **Bazaar**'s [9] main focus is simple and easy user interface; though internally using elaborate unique revision identifiers, it presents primarily the RCS-like revision numbers to the user, and its user interface has direct support for some common workflows like emulation of the centralized version control setup *(checkouts)* or **Patch Queue Manager**, an advanced alternative to all developers pushing to a central repository.[24] Bazaar is also written in Python, making the same tradeoffs as Mercurial.

3.7 Darcs

Darcs [19] is a somewhat unusual specimen in the arena of version control as it focuses on changes as the first-class objects instead of trees; any particular version of the repository is merely a combination of patches on an empty tree, and manipulating the patches is the primary focus.

Darcs supports single branch per repository and names the patches based on user input; particular combinations of patches can be tagged to mark specific versions of the tree.[25] Darcs does not organize the patches in a particular graph, they are structured more like a stack, though a graph could be imagined based on interdependencies between the patches; merges of several patches create auxiliary patches representing the merges.

A particularly interesting feature of Darcs is precisely how it handles patch merges — it uses a formal system called *patch algebra* (6.6) to describe dependencies between patches and basic operations on the patch stack required to properly merge two sets of patches (branches).

Darcs is written in Haskell, which is somewhat exotic choice, but is well portable. Unfortunately, as far as we are aware it suffers from severe performance problems (as discussed in 2.3) when used on large repositories with long histories; more efficient merging algorithms as well as proving formal correctness of the underlying patch theory is subject of ongoing research.

4 Storage Models

The storage model of a system can be quite different from the logical design in order to accomodate for practical considerations of space and time effectivity. Here we shall explore possible approaches to permanent storage of the data and metadata tracked by version control. All version control systems differ in implementation details and various individual tweaks; we will not try to be exhaustive here and will focus only on the most widespread and most interesting ideas.

First, we will dwell on the currently popular algorithms to compare two objects and create a *delta*: complete technical description of their differences.[26] Second, we shall elaborate on possible ways of representing an individual delta as well as whole single-file history, then we will look at the approaches on organization of the whole repository.

4.1 Delta Algorithms

There is a certain variety of delta algorithms in use, and there can even be multiple algorithms used within a single system — some algorithms provide suitable and sensible output for human review while others are focused on finding the smallest set of differences and providing minimal required output for storage.

4.1.1 Common Techniques

First, we shall describe several common techniques used when dealing with deltas. The

[24] Developers send their pull requests or actual patches to a public mailing list, where they can be acknowledged or vetoed by others and then a robot automatically picks them up, checks if they pass the required testsuites and eventually merges them to the main branch. [47]

[25] In fact, a Darcs tag is itself a patch that is however empty and depends on all patches currently applied.

[26] E.g. a series of insert–copy commands or a "diff" tool output text.

first technique worth mentioning is used to our knowledge by all commonly used systems newer than CVS: *delta combination* [45] [68]. When applying a chain of deltas, they are first merged to a single large delta, and only then applied. This reduces the number of data copies from $O(deltas)$ to $O(1)$ and takes advantage of cache locality, dramatically increasing performance of applying delta chains.

Another interesting technique used most notably in Subversion is using *skip-deltas* when storing delta chains. [45] In order to limit chain lengths, the deltas aren't always created against neighboring revision but instead construct a binary tree alike structure — thus, reconstructing particular revision does not require $O(revs)$ deltas but only $O(\log revs)$. However, when adding new revisions, parts of the tree or the whole tree needs to be reconstructed, and the deltas are usually larger. Newer systems usually solve the problem of too long delta chains by simply inserting the whole revision text when the chain becomes too long or large in size.

4.1.2 Myers' Longest Common Subsequence

Perhaps the most widely used algorithm was proposed by Eugene Myers in [7] and is based on recursively finding the longest subsequence of common lines in the list of lines of the compared objects; it is used most notably by the **GNU diff** tool [33], but also as the user-visible diff algorithm of most systems. The main idea of the algorithm is to perform two breadth-first searches through the space of possible editations (line adds/removals/keeps), starting from the beginning and end of the object in parallel; when the two searches meet, we have found the shortest editation sequence. The algorithm has quadratic computational complexity, so it is unsuitable for very large inputs.

As noted in [68], while this algorithm is quite suitable for human consumption and does produce optimal longest common subsequences, it does not produce optimal binary deltas since it weights additions, changes and removals all the same, while binary deltas need to actually record only added data and removed data does not need to be spelled out.[27]

[27] However, this should be easy to work around by adjusting the breadth-first search to add 0-length edges at the beginning of the queue and 1-length edges at the end.

4.1.3 Patience Diff

An interesting variation of the above algorithm was recently introduced as the default diff method to the Bazaar system, devised by Bram Cohen. [10] A common problem with using the Myers' algorithm is that the common subsections search misaligns additions of parentheses on separate lines or other simple syntactic constructs, making it harder to view by humans.

The main idea in this algorithm is to "take the longest common subsequence on lines which occur exactly once on both sides, then recurse between lines which got matched on that pass." [61] Also, it achieves better computational complexity by using Patience Sorting [36] to find the longest common subsequence.[28]

4.1.4 BDiff

The bdiff algorithm [68] improving the Ratcliff–Obershelp pattern recognition [48] comes from the Python **difflib** library [53] and is used by Mercurial, both for delta storage and providing difference listing to the user. In contrast to looking for the longest common subsequence, it searches for the longest common continuous substring within the objects and recursively in the parts preceding/succeeding it. It is quadratic like the Myers algorithm, but has been measured [68] to have better average performance and output more human-friendly deltas.

4.1.5 XDelta

The tool **xdelta** [72] [73] popularized use of the Rabin Fingerprinting technique [23] for delta generation: it is effectively a simple modification of the Rabin–Karp substring matching algorithm, inspired by the work on **rsync** [22]. Subversion [60], Monotone [41] and Git [28] use this variation for internal storage — generating one-way binary diffs between arbitrary non-textual blobs.

The algorithm produces COPY and INSERT instructions on the output and depends on fast hashing. The first object is divided into small blocks[29] and each block is hashed and added to

[28] The referenced paper describes finding the longest increasing subsequence. To find longest common subsequence, a sequence of lines corresponding to one object is created, with the elements being line numbers of the same lines in the other object; lines that are present more than once or never are ignored in the search.

[29] The window size is 16 bytes in Git; small power of 2 is recommended in general.

a hash table (proportional to input size). Then, a running hash of the same-sized window is being computed on the second file; the moment a match is hit, all found blocks of the first object are stretched to largest possible matches to the current position in the second object and the COPY instruction is generated for the largest match. INSERT instruction is generated for unmatched data in second object.

This algorithm is very fast (having much better computational complexity compared to Myers' — $O(n)$ vs $O(n^2)$) and produces efficient diffs for arbitrary data, however it is entirely unsuitable for human consumption.

A similar scheme is also used in zdelta [74] with reportedly better efficiency and with code released under a more permissive licence. ZDelta reuses the LZ77 algorithm [4] for finding the matching parts and Huffman codes for encoding the result.

4.2 Delta Formats

4.2.1 Trivial Approach

Not really a *delta* format by itself, the most naive approach to store history is to simply store every version of every object separately, possibly using some compression method (not taking advantage of historical context of the object). This is the obvious format used for "poor man's version control" methods like archiving compressed snapshots of the whole project at regular intervals and backing them up. However, perhaps surprisingly this method does see usage in some modern systems as well — namely, it is used as the basic in-flight storage format in Git.

4.2.2 Unified Diff

The unified diff format [69] — albeit not used in the changes storage itself — is the de-facto standard for interchange of plaintext changes, being the usual format used in "patch" files.[30] For each file, the diff contains all the *change chunks* — the areas of the file where changes occured. Each chunk consists of the chunk header localizing the chunk in both versions of the file (starting line number and line length of the chunk), *context* lines immediately surrounding the changed lines (preceded by a space) and the change itself: added lines preceded by a `+` sign and removed lines by a `-` sign; modified lines are written as a `+`/`-` line pair. Unidentified lines in the diff are ignored, which allows version control systems to include custom metadata in the diff.[31]

This format has number of advantages: it is easy to generate, it is easy to inspect by a human for the purpose of change review (including easily tweaking details of the change in the unified diff itself), it is easy to apply automatically[32] and very importantly, it is possible to apply the differences to another version than the one the diff was created from. This is allowed by the presence of the context lines: even if the starting line numbers in the chunk header do not fit the context, the program can automatically search the vincinity of the area for the context lines and apply the patch fuzzily. At the same time, the context lines provide a way to verify that the change is still applicable as-is to the target version.

We should also mention a Git extension to the unified diff for showing diff representation of merges — *combined diff* [26] (see figure 1 for an example). Instead of a single column with the ± marks, multiple columns are included, one per merge parent, allowing concise description of which lines differ against which parents. Furthermore, by excluding hunks that contain only changes coming from one of the parents, we get a diff describing only changes actually introduced by the merge itself — manual resolutions of merge conflicts.[33]

4.2.3 VCDIFF (RFC3284)

The VCDIFF format [64] deserves a brief mention, being the official standardized way to exchange binary deltas over the network, being specified by RFC3284. However, to our knowledge no widely used version control system actually uses this format for anything, though it is the default output generated by the xdelta tool (4.1.5). The main reason is that it is easier and faster for individual systems to use their native delta formats even for data interchange and that the VCDIFF format is rather baroque,

[30] Sometimes, a very similar format called "context diff" is also used, however unified diffs are much more common.

[31] Commonly, informational-only notes about the compared revisions are inserted. However, e.g. Git uses this to extend the diffs with file renames/copies information and can make use of the metadata when applying the patches.

[32] UNIX provides a standard `patch(1)` tool. [52]

[33] Conflicts are further explained in 6.2.

```
diff --cc git-am.sh
index 75886a8,4dce87b..b48096e
--- a/git-am.sh
+++ b/git-am.sh
@@@ -9,9 -8,9 +9,9 @@@ git-am [options] <mbox>|<Maildir>..
  git-am [options] --skip
 -d,dotest=       use <dir> and not .dotest
 +d,dotest=       (removed -- do not use)
  i,interactive   run interactively
- b,binary        pass --allo-binary-replacement to git-apply
+ b,binary        pass --allow-binary-replacement to git-apply
  3,3way          allow fall back on 3way merging if needed
```

Figure 1: Combined diff example. [26] The `diff` and `index` lines are Git-inserted metadata, and the trailing text at the `@@@` line is only to give the reader a semantic context of the hunk.

implementing it is not trivial and it is generally not as compact as the native protocols.

4.2.4 RCS Delta

RCS uses the simplest representation of changes sequence that also spilled over to CVS and heavily influenced the SVN BerkeleyDB schema.[34] [54] The representation is in text format and is optimized for quick access to the latest revision, which is stored at the top of the file and always in full. Older revisions are then stored each as a delta against the succeeding one. In case of branches, the delta direction is reversed and newer revisions are represented by forward delta — thus, to get a tip of a non-trunk branch, delta sequence all the way from the latest trunk revision to branch base and then back to the branch tip must be applied.

The delta itself is simply composed by lines each containing an add a or delete d command, followed by line number in base data and number of lines, with the lines to insert by a following the command immediately.

Aside of very slow reconstruction of non-trunk branch tips and old revisions in general, the major issue with this format is the requirement to rewrite the whole file in case of commit, making the repository database much prone to lock contention and data corruption and making creation of new revision quite an expensive operation.

[34] SVN FSFS format is quite different and perhaps most resembles the Mercurial revlogs (4.3.3). We chose not to describe SVN formats in detail since we deem them not particularly interesting in themselves.

4.2.5 Weave

The weave format introduced by SCCS [63] used to be rather obscure for long time, but currently many version control developers are familiar with it due to its prominent usage by BitKeeper and some newly developed merging algorithms. Still, it has not seen widespread adoption.[35]

Instead of storing each revision separately, this format intersperses all revisions, listing all the lines that ever appeared in the file together with the list of revisions they appear in. Then, retrieving any revision means simply extracting all the lines that have the revision in their set.

This format shares many disadvantages with the RCS delta format, especially with regard to locking and data corruption potential. In addition, retrieving any revision is uniformly proportional to history size. On the other hand, the line annotation[36] is available very cheaply (compared to systems using delta-based storage where retrieving this involves walking all the history) and having per-line history information enables some interesting merging methods to be used (see 6.4).

[35] The Bazaar project actually used it for brief period but then switched to a more traditional delta chain representation for the benefits of append-only data structure. [13]

[36] Information about who, when and in which revision added a particular line.

4.3 Repository Database Formats

4.3.1 Trivial Approach

The trivial approach is to simply use separate files for appropriately fine-grained objects. This usage is very common in many systems; it is the approach taken by CVS and Mercurial where every file in the project directory has its history stored in corresponding file in the repository. It is also the basic storage format of Git, where each object (commit, tree or blob — particular state of a version-tracked file) is stored in a spearate file.

4.3.2 Git Packs

As mentioned above, Git simply stores its first-class objects directly on the disk as its primary format. However, as explained earlier this is very inefficient for long-term storage; for that purpose, Git Packs offer a storage format showing extreme size effectivity in practice.[37]

A Git Pack [26] consists of a pack file and index file. The index file provides a simple two-level index of the objects keyed by the object id. The pack file contains the compressed Git objects of all basic types — `commit`, `tree` and `blob` — but can also accomodate an extra object type `delta`:

The pack-specific delta objects represent a given object by describing binary differences against another object of the same type,[38] as yielded by the xdelta algorithm (see 4.1.5). In theory, any object of the same type will do as a base, but Git sorts the objects using a special heuristics and then considers n following objects in the list as delta candidates,[39] picking the smallest delta generated.

The list heuristics is the core of the Git Pack efficiency. [16] [28] The important criteria are that blob objects are sorted primarily by the filename they were first reached though, with the last characters being the most significant ones. (Thus, all `.c` files are grouped together, then all files with the same name in different directories, etc.) The secondary sorting order is by object size, largest objects first — thus, deltas are generated from larger objects to smaller ones, allowing more efficient representation.

Note that the delta order is in almost no way related to physical ordering of the objects in the pack file. There are no theoretical restrictions on that, however Git orders them based on breadth-first search on the graph of all objects starting from the top commit objects on the defined branches. In other words, the objects necessary for recreating the latest commits are lumped together at the beginning of the pack file. This particular ordering greatly optimizes the I/O operations locality for the most common I/O patterns.

One specific exception in the physical ordering is that in case of delta objects, the base object is stored before the delta object; this ensures locality when reconstructing delta chains. In practice, this rule does not reorder the objects significantly — as the *Linus' Law* states: "Files grow" [16]; thus, the base objects (larger files) tend to come before delta objects (smaller files) anyway and while it is not enforced, the delta objects tend to become "backward deltas" most of the time.

Git Packs use no locking mechanism — they are not created on the fly, but only at certain points in time during a "garbage collection" operation over the repository, from loose one-per-object files created in the repository by regular usage.[40]

4.3.3 Mercurial Revlogs

Another popular data structure are so-called revlogs used by Mercurial. [68] [38] While Git packs are designed around heuristics, guarantee very little about worst-case performance and are optimized for work with hot cache,[41] revlogs can guarantee good worst-case seek performance (at the cost of less flexibility for improving average performance) and are tailored for cold cache access.[42] Unfortunately, while there are some casual benchmarks available on the web, detailed comparison with clearly defined cache status to confirm these expectations is yet to be published.

[37] Git also uses the same format for native network data transfer.

[38] Or, of course, another delta object. However, the delta chain is upper-bounded, by default by the value of 10.

[39] By default, the delta window is $n = 10$, but commonly, values like 50 and rarely even 100 are used for generating ultra-dense packfiles.

[40] One exception is when new commits are pulled by the native protocol, which itself is in the Git Pack format; the data is then saved directly as a pack.

[41] That is, when the user uses the system continuously for longer time and most of the relevant data is cached in memory.

[42] With "cold cache", most of the data are still on disk.

Mercurial uses separate revlogs for each tracked file and two extra revlogs for the manifests and changesets. Each revlog consists of an index file for revision lookup[43] and data file containing a hunk per revision, every hunk being either a full revision in verbatim or a delta against the previous revision — the total size of the delta chain is limited by a small multiple of the revision size.[44] In case of the revision being represented by a delta, the index will contain pointer to the base revision so that all the necessary data can be read from the file in a single swipe. Revisions from different branches are all stored in a single revlog, but when forming delta chains only the order of recording the revisions matters, not their branch.

When recording new revision in a revlog, the revlog is locked for writing and a new hunk is appended to the data file and then recorded in the index. Since the revlogs are append-only, no locking is required for readers and it is a matter of simply truncating the file to recover from an interrupted write. Furthermore, when working with the whole tree (both reading and writing), all the revlog files are visited in a fixed order compatible with the standard system files ordering; this way, the filesystem can maintain the most desired on-disk layout and using standard system tools to copy a repository will result in a good on-disk layout as well.

5 User Interaction

In the development of new generation version control tools, the importance of good user interface has been underestimated in the past, while it tends to be one of the most important factors for users when deciding which system to use; good user interface is therefore one of main considerations in the development of currently widely used tools.

Traditionally, all of the tools provide command-line interface, which tends to be the main way of interaction with the system from the side of the developers. However, many authors also explore ways to interact with the system efficiently using GUI, both for actually operating the system and merely exploring the history. Another separate area are web interfaces for inspection of repositories.

5.1 Command-line Interfaces

One important conisderation with the command-line interfaces is to keep in mind that users are actually going to *type* many commands — e.g., GNU Arch requires usage of extremely long revision identifiers, and for systems using hashes for object identification, it is important to allow shortcutting mentioned earlier.

Another considerable point is to design the interface carefully and logically around a minimal set of commands, also hiding implementation details whenever possible (e.g., by trying to intelligently default to a suitable merge algorithm instead of requiring the user to decide on one). Git provides currently about 140 commands, while large portion of them is in fact almost never interesting to end users and is useful only as helper commands within scripts or other commands.[45]

Note that even regarding traditional command-line interfaces, there can be interesting variations — for example, Git extends the basic edit–commit workflow with the optional *Git index* concept, presenting a staging area inbetween the working tree and next commit; see 3.6 for details.

User interface can take very unusual forms; a specific one is integration of version control within the filesystem interface. In that case, accessing a particular revision of a file is usually performed by simply opening a specifically crafted interface. This concept goes as far back as to the **VMS** operating system where the filesystem interface had builtin support for file versioning. [71] A configuration management system with version control capabilities **Vesta** provides primarily filesystem-based access to individual revisions (furthermore with $O(1)$ access guarantees). [65] There also exist alternative filesystem-based interfaces for more traditional version control systems as well (though

[43] A slightly modified RevlogNG format used by newer Mercurial versions supports interleaving of index information and data hunks.

[44] Mercurial 1.0 uses the value of 2.

[45] To simplify the usage of Git, we created a frontend Cogito that used lowlevel Git interface to create an easy-to-use version control tool carefully designed to simplify many workflows and present a simple and consistent set of commands to the user, with consideration for users used to older systems like CVS or Subversion. Since most features provided by Cogito have been adopted or obsoleted by upstream Git development, Cogito is not maintained anymore (but we still consider the Git UI somewhat lacking).

mostly in early development stages), e.g. **BzrFS** [12] or **GitFS** [31].

5.2 Basic GUI Tools

There are several types of GUI tools. The most obvious candidates are all-in tools covering most of the system functionality, often reusing other helper applications (such as history browsers) — most systems have tools like that, e.g. **WinCVS**, **git-gui**, **Bzr-Gtk**, etc.

An interesting approach is to integrate the version control system with basic system shells or file managers (like the Windows Explorer or Konqueror); version-controlled files include the state in their icon (clean, modified, containing merge conflicts, ...) and all versioning operations are available through the context menu. The classical example being **TortoiseCVS**, clones like **TortoiseSVN**, **TortoiseBzr**, **TortoiseGit**, **TortoiseHg**, etc. also exist. Similarly, version control capabilities are often integrated with IDEs like Eclipse.

Another family of tools are simple GUI helpers designed merely to be used in concert with the classical commandline controls; this has been mostly pioneered by BitKeeper, providing citool, mergetool etc. to allow to perform operations like reviewing to-be-committed state or merging files more comfortably. Currently, e.g. Git frontends like git-gui or tig provide a way to control to-be-committed content down to the per-hunk level of diffs and there exist system-agnostic tools for graphical review of diffs (**xxdiff**, **kdiff3**, ...) and even merging assistance (especially **meld**).

Visual UI tools do not need to be always graphical. Many developers prefer exclusively terminal work and sometimes it may be necessary to work on a remote server over ssh session; **tig** is a good example of terminal tool that provides ASCII-art[46] history diagrams and frontends many of Git's functionality (including index management down to diff-hunk level), without requiring any graphics rendering.

5.3 History Browsing

With the advent of distributed version control, it became much more important to properly visualize the project history — suddenly it is not a simple tree of commits, but a more complex graph with possibly many branchings and merges and the relationships between revisions may not be immediately obvious to the user.

The first popular history visualizer were BitKeeper's revtool and **Monotone-viz** [42] which focuses on visualising the revision graph, offering it as the main visual object to the user; clicking on nodes will display revision details. During Git infancy, the tool was adopted as Git-viz, but it did not see widespread usage; however, similar tool is being used in Bazaar as part of BzrTools.

Instead, a much denser history presentation gained popularity in the Git community within a tool **gitk**. Its primary visual objects are the revisions themselves, with the first line of the commit message (traditionally regarded as kind of "subject") and authorship information shown for each commit; the graph is still shown, however only in narrow area at the margin of the screen, with nodes aligned to the line structure of commits. Reused in other Git user interfaces (both graphical ones and tig), this way of visualization became popular in the Mercurial community as well (**hgk**), Monotone even produces similar kind of output in ASCII art with its native `mtn log` command; there also exist projects like **git-browser** which use AJAX technologies to bring this interface to the web.

5.4 Web Interfaces

Important part of user interface of a version control system is a web interface — most projects have some kind of web interface available so that people wanting to look at the source and the history do not have to install the particular version control tool and download the whole project. Good interface also should not assume users' familiarity with the particular system, since it gets much more diverse audience than other tools of the system.

Historically, **cvsweb**, **viewcvs** and similar tools for SVN focused on primarily presenting the tree of files, with support for inspecting individual history of files when needed. While similar approaches are also present in some web interfaces for distributed version control systems (e.g. **GitHub**), the paradigm shift to primarily presenting the project history is visible here as well, in the widely used **gitweb** interface and the **hgweb** clone. On the project front page, the user is not presented with root directory of the file tree, but instead with the summary of

[46] Graphics constructed in text mode from standard letters and symbols.

the latest commits, available branches and tags, etc. — browsing the tree is of course possible, but not the primary activity anymore.

Something of a rarity, the **Gist** service[47] uses modern web technologies to allow users build whole repositories from scratch from within the web browser; the user is presented by a text-area where they can paste content (akin to various *pastebin* services), but further modifications are then stored in git revision history, it is possible to add multiple files and eventually pull the history. It is unclear to us how often it is actually used as anything more than a trivial one-off pastebin, however.

6 Merging Methods

One of the major past challenges with systems encouraging massive branching was devising a powerful merging mechanism that would make the merging process sufficiently easy, quick and robust for frequent usage. This ability is crucial, otherwise people will become reluctant to create branches and the main strength of distributed version control will dwindle. On the other hand, practice shows that merge algorithms should be simple enough to be easily predictable, even if the tradeoff would be to require manual interventions by the user more frequently — common problem of algorithms not based on three-way merge is that while they may generate less conflicts, their merge resolutions may be relatively unintuitive and some of their properties are controversial and unfamiliar to the users.

In this section, we will describe the most common merge strategies (*three-way merge* and *recursive merge*) as well as some more advanced theoretical developments (*mark-merge*, *PseudoCDV merge*). We will also consider the *patch algebra* based way of merging in changeset-oriented systems.

There is a great number of considerations when designing and comparing merging algorithms and we cannot cover them all here; we recommend the pages dedicated to merging algorithms in the RevCtrl wiki [55] for detailed study of all the aspects.

6.1 Fast-forward Merge

First, we should notice what happens in the revision DAG when a "blank" merge is actually performed: consider merging revisions 1.2 and 2.4 as shown in the figure 2. On the branch 1, no new revisions appeared since its last merge with 2.2, thus merging the 2.4 revision will effectively merely copy this revision over to the branch 1. A system might create a merge-node 1.3 anyway, but this may not be desirable since many useless merge-nodes will be created in the extremely common scenario of users only tracking some project in read-only fashion: every time they update their copy, no local revisions will be found locally but the merge will create a merge-node revision to represent the update anyway.

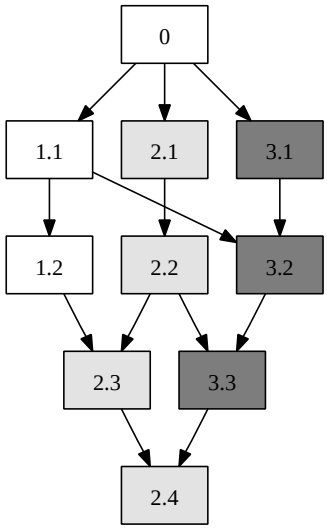

Figure 2: A history DAG example

To avoid this problem, a "virtual" merge strategy of *fast-forwarding* is used; in case one of the merged branches contains no new revisions since the last merge, it is simply repointed to the head commit of the other branch. Thus, in our scenario, branch 1 is simply repointed to the revision 2.4! Then, when committing a new revision to branch 1, 2.4 will be treated as the new fork point.

This technique allows branches of the DAG to converge whenever possible, however it also breaks commit parents ordering — after fast-forwarding to a different branch, merges with the original branch will appear "flipped". The users either cannot rely on the parents ordering, or they have to constrain the system usage by a policy to disallow scenarios that would provoke a fast-forward.[48]

[47] Part of the GitHub cloud.

[48] Fast-forwarding can frequently occur in tight push-pull loops performed concurrently on a shared central

6.2 Three-way Merge

Three-way merge[49] is the most common method of merging in use. Assume merging two version nodes x and y into merged version m. Let b be the *merge base* — the "nearest common ancestor" of x and y. Then, using difference function[50] $\Delta\colon (text1, text2) \to delta$ (and delta applying function $\Delta^{-1}\colon (text1, delta) \to text2$):

$$d_x = \Delta(b, x)$$
$$d_y = \Delta(b, y)$$
$$d = d_x \cup d_y$$
$$m = \Delta^{-1}(b, d)$$

Simply the combination of differences between base and both merged versions are applied to the base version again and the result is the merged version. The problematic part here is determining d — if d_x and d_y concern discrete portions of the base version, combining them is trivial.

However, when both change the same part of base version, each in a different way, they cannot be combined in a simple way. A *merge conflict* arises, only the non-conflicting parts of d are applied and for the rest the user is required to choose between $m_x = \Delta^{-1}(b, d_x)$ and $m_y = \Delta^{-1}(b, d_y)$ or combine m_x and m_y manually — usually the system inserts both variants to the file, separated by visible one-line markers.[51]

In general, in single-shot usage[52] the algorithm can work on any b merge base version. However, the choice of a particular b has major impact on the size of d_x and d_y, affecting the conflicting portion of d. Thus, b should be chosen intelligently, usually by taking the *least common ancestor* (LCA).

In case of a tree, LCA can be simply described as the node nearest to the tree root on a non-oriented path from x to y, and it will be always unique as follows from tree properties. However, this approach is satisfying only for the first-time merge of the branches of x and y — if a merge between the branches is performed repeatedly, the same b will always get chosen, even though the previous merge outcomes would be much more reasonable choices as the size of d_x and d_y can be assumed to generally increase over time.

This is the main motivation behind making the history a DAG instead of a tree; when recording the m version, both x and y are marked as its ancestors instead of just one of them. When determining the LCA[53] of descendands of x and y later, m can then be used instead of the original b. Thus, for example for the graph in figure 2 earlier, the base of the 2.4 merge of 2.3 and 3.3 would be 2.2.

However, in this case we lose the guarantee of b uniqueness, since multiple discrete paths between x and y can exist: to deal with this problem, recursive merge extension of the algorithm has been devised.

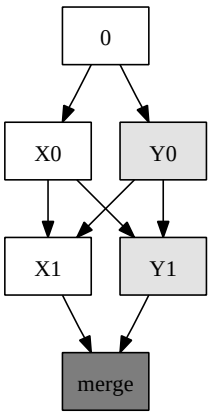

Figure 3: A criss-cross history graph example

6.3 Recursive Merge

Consider the *criss-cross merge* scenario [56] shown in figure 3. In this case, b can take the

repository. One way to avoid them is to tell the developers not to merge the changes their peers have done in parallel and pushed earlier, but instead to *rebase* their own changes on top of the already-pushed ones: this operation rewrites the history locally, sequencing two commits performed in parallel.

[49] Unfortunately, we were not able to find out who introduced the technique of three-way merging, currently commonly known and widely used; the earliest reference we were able to find is in the RCS paper. [54]

[50] See 4.1 for the tour over various functions.

[51] In some cases, the system can still resolve simple conflicts, for example if they concern only whitespace characters or perhaps even when merging a simple reformatting change with a semantic change. Putting more inteligence into conflict resolution is one of the active areas of merge research. However, silently mis-merging incompatible changes can have dangerous consequences, so making the resolution more intelligent at the risk of increasing error rate can be very harmful.

[52] When the algorithm is applied on two arbitrary versions without historical context. In a generic DAG, selecting inappropriate b within repeated merges can lead to information loss, as described below.

[53] In case of a DAG, we prune paths that are supersets of other paths from our consideration.

values of both X_0 and Y_0, however taking any of these can lose data from the other branch depending on the nature of the merge — consider conflicting change in X_0 and Y_0, in X_1 resolved by taking X_0 version exclusively and in Y_1 by taking Y_0. Then, for $b = X_0$, Y_0 version will be cleanly merged and vice versa — clearly, conflict should be generated instead.

This particular problem and its more complicated variations led some researchers to abandon the three-way merge altogether and focus on other merge methods instead, some of them described below. However, the problem can be at least partially solved within the context of three-way merge as well, though the practical impact has to be yet evaluated carefully.

A basic rule should be that the system, detecting multiple possible merge bases, does never randomly select one but instead asks the user for intervention, chooses a proper merge base or possibly intelligently constructs one.

One naive approach we devised is to take $\mathrm{LCA}(X_1, Y_1)$ recursively and repeat until a single revision comes out. The practicality of this approach is dependent on the development strategy adopted by a particular project — if the criss-cross merge is not repeated for too long and a common base for both branches appears once in a while, the LCA will stay in a reasonable distance from the merged versions and conflicts would generally stay within reasonable scale. However, for long-term criss-cross merges this approach degenerates to something similar to a tree-based history, and it still leaves some more advanced criss-cross problems unsolved [1].[54]

An important enhancement implemented by Fredrik N. Kuivinen is to actually perform the merges themselves recursively: [27] [56]

$$
\begin{aligned}
|B| = 1: &\quad b(B) = B_0 \\
|B| = 2: &\quad b(B) = M(\mathrm{LCA}(B_0, B_1), B_0, B_1) \\
M(B, x, y) &= \Delta^{-1}(b(B), x \cup y) \\
m(x, y) &= M(\mathrm{LCA}(x, y), x, y)
\end{aligned}
$$

That is, in the example above:

$$m = \Delta^{-1}(\Delta^{-1}(0, X_0 \cup Y_0), X_1 \cup Y_1)$$

It is easily visible that this works properly in case of no conflicts, and reduces conflict rate by working on much finer grain level. In case of conflict, the main idea of the algorithm is to simply *leave the conflict markers in place* when using the result as a base for further merges. This means that earlier conflicts are properly propagated as well as conflicting changes in newer revisions. However, some rare edge cases are still mishandled as described in [56].

6.4 Precise Codeville Merge

A weave-based merge algorithm *Precise Codeville Merge* [62] has been devised by the **Codeville** version control system.[55] It later turned out that Codeville probably independently invented a merging method very similar to what BitKeeper uses internally.

Instead of comparing the to-be-merged revisions to a base revision, the algorithm directly performs a two-way merge between the two revisions and then uses the weave-embedded history information to decide which side of each conflicting hunk is to be chosen.

Consider $v_i(r)$ being the number of state changes (addition or removal; also called *generation count*) of line i at revision r. For each revision, have a list of line chunks: l_x, l_y. Then, the weave of the file is iterated line by line:

- if a line is not in either of the revisions, nothing is done
- if a line is found to be only in one of the revisions, it is appended to appropriate line chunk list l_x or l_y and revision precedence flag p_x is set if $v_i(x) > v_i(y)$, or p_y if $v_i(y) > v_i(x)$ (obviously, $v_i(x) \neq v_i(y)$)
- if a line is found to be in both revisions, either:
  1. p_x and p_y are unset and this is part of common lines block and appended to the output
  2. only p_x or p_y is set and this is non-conflicting change (then, l_x or l_y respectively is appended)
  3. both are set and a conflict block from l_x and l_y chunks is generated; in the latter two cases, $l_x \leftarrow l_y \leftarrow \emptyset$ and $p_x \leftarrow p_y \leftarrow 0$.

[54] It must be also considered that if the merge generates unnecessarily large conflicts, risk of human-induced mismerge raises considerably.

[55] The system is not otherwise covered here since it has never seen wider usage and is not developed anymore; it can be probably said that its main purpose ended up to be to research various merging algorithms.

The advantage of this algorithm is the absence of the need to walk the revision graph at all (provided that the system uses the right storage format, which can be difficult while keeping the desirable append-only properties of the most common ones), simplicity and handling of most of the problematic edge cases of the three-way merges. Also, this algorithm handles cherry-picking naturally well compared to three-way merge.

The disadvantage is a requirement for the weave structure potentially expensive to emulate if the system uses different storage method,[56] and somewhat controversial usage of the v_i criterion — it is not immediately obvious that lines with more state changes should always win against these with less. Due to the fundamental difference in operation to the much more widely used three-way merge, in cases it chooses different result its operation may be confusing to the users.

6.5 Mark Merge

The mark merge (or *-merge) algorithm [35] is somewhat unusual. First, all the other merge techniques described here are designed to merge vectors (file content), while this one is a *scalar merge* algorithm (for example for directory items — the executability flag or file name when merging renames). Second, mark merge has a precisely specified user model and formal proof that the algorithm matches the user model.

The aim of mark merge is to define the most sensible way of merging scalar values while giving foremost priority to explicit user choices. First, a *mark* means explicit decision by the user on the scalar value; marked revision is such that user explicitly set the value in this revision. Let us have version r, then $*(r)$ shall be set of minimal marked ancestors of r — that is, taking graph of r ancestry and reducing it to marked revisions, set of revisions with no descendants. Then, when merging versions x and y with values v_x and v_y:

- if $v_x = v_y$, the value shall be also the merge result
- if all $(x)*$ revisions are ancestors of y, v_y shall be merge result
- if all $(y)*$ revisions are ancestors of x, v_x shall be merge result
- otherwise, throw a conflict

For detailed formal user model and proof of correctness please refer to [35].

6.6 Patch Algebra

In Darcs (3.7), a particularly elegant formal system *patch algebra* has been developed for merging two sets of patches in a changeset-oriented version control system. [20]

The main idea of patch algebra is to provide natural and formally-backed semantics for merging within such a system.[57] All patches are defined to have a set of dependencies on other patches[58], inversions, and independent patches are allowed to *commute* on the stack. When merging two sets, independent patches are appropriately commuted (sometimes requiring use of inversion patches to work around fuzzy applications) and for conflicting patches, user resolution is requested; the merge operation is then recorded in a special merger patch.[59]

Given the way the to-be-merged merger patches are unwound, it is believed that the merge behaviour with regard to the pathological cases like criss-cross merge are similar to the recursive merge algorithm or better, since Darcs can make use even of intermediate changes between the two merged revisions and their base that the three-way merge approach cannot see. [56] [8] However, we are not aware of any work formally analysing and rigorously comparing the properties of Darcs merges with the other approaches.

7 Future Research

On the technical side, some of the existing systems are not portable enough to be usable

[56] Vesta currently uses this merge method and constructs weaves on-demand; scalability of this method was not studied, however. [51]

[57] But please note that while Darcs makes an attempt for formal description, it is still lacking in many aspects both regarding precise definitions and many proofs missing. We must however mention a relatively little-known theoretical work [2] inspired by Darcs formalism and attempting to rebuild it on sound foundations.

[58] The dependencies are either logical as explicitly specified by the user, or syntactic — a patch is depending on all patches lines of which it touches — as autodetected by the system.

[59] The actual mechanisms used to preserve all the desirable properties of the patch stack are rather complicated — please refer to [20] for full description.

equally well in different operating systems (especially with regard to Microsoft Windows vs. Linux and other UNIX systems) and the user interfaces are still lacking in many aspects, as well as documentation.

Regarding the theoretical side, we believe three main areas need research most urgently — content movement tracking, management of third party changes and more precise formal specification.

7.1 Fine-grained Content Movement Tracking

More fine-grained tracking of changes is required to improve merging algorithms, accounting for changes and following history of arbitrary parts of content; as noted in 2.5, some of the systems provide way to follow file history across renames, **Git** introduces the *pickaxe* mechanism, but better infrastructure is required for practical usage especially during merging and allowing visual representation and usage.

7.2 Third-party Changes Tracking

The current distributed architectures still are not flexible enough to accomodate for fully distributed development in open-source projects. When a third-party developer publishes some changes and the upstream developers want to merge it only with certain reservations, the third-party developer can either do the required changes on top of his previous changes, leaving the history in a certain degree of disarray and making his proposed chances increasingly harder to review, or rewrite the history of his changes, which creates problems for others who have already accessed the published version since the rewrite essentially created a fresh new branch unrelated to the older one.

Similar problems arise when the third-party changes need to be ported to new upstream version, or when a distributor wants to maintain series of changes in their version of the product; using the traditional merge mechanisms, it quickly becomes difficult to track the changes against upstream cleanly as the ability to generate a diff of single change against latest upstream version can be crucial.

Thus, many people make use of the distributed nature by only maintaining the changes locally and still using patches for changes interchange; or, they use the version control system for tracking upstream and some special frontend on top of that[60] for tracking their changes in form of stack — this has user interface issues and it is problematic to publish the repositories since the history is changing constantly. Another alternative (often used by open-source software distributors when maintaining their packages) is to version-control the patches in their diff form directly.

Using patches for changes interchange can have some benefits with regard e.g. to code review practices, but frequently it is merely extra overhead and it incurs extra load on the third-party developers. A new design that would allow seamless external changes tracking needs to be developed.

7.2.1 TopGit

We have created **TopGit** [66], a layer over Git that aims to solve the task of maintaining third-party patches properly, allowing for fully distributed development and proper history tracking of all the changes.

We start off from the concept of *topic branches*, popular Git technique of having a specialized branch for development of each independent change, then merging them all together; this is merely a "design pattern", needing no special tool support. However, we extend this in two ways:

1. We create a directed acyclic graph from the topic branches,[61] giving each branch a list of dependencies (other branches).

2. We store metadata for each topic branch within its file tree: the branch description and authorship information within `/.topmsg` file and the list of branches it depends on inside `/.topdeps`.[62]

The topic branches can depend on other topic branches, but also on regular Git branches. Typically, in the main set of third-party patches, each patch will have its own TopGit branch depending on a Git branch tracking the upstream development, then possibly some third-party patches will have dependencies on some of the TopGit branches if they further extend

[60] E.g. **mq** for Mercurial, **StGIT** for Git

[61] Thus, conceptually one level higher than the DAG of the *commits*.

[62] These two files are excluded from merges of dependencies.

the third-party changes, and at the top of the graph will be usually at least one so-called *staging branch* that introduces no changes on its own (as opposed to *content branches*) but ties all the third-party patches together by listing them all in its dependencies.

TopGit provides tools for creating topic branches, querying them, pulling topic branch updates and merging them with local work, but most importantly a command for updating the branches, which recursively updates all the dependencies of a branch, then merges them in.

The project reached basic practical usability and at the time of this writing is used for maintenance of some Debian packages, but its user interface leaves a lot to be desired and some of the functionality is not fully implemented.

Especially problematic is the task of *removing* a dependency of a patch in a way that preserves the history, makes it possible to easily re-add the dependency later and does not spread the dependency removal to further depending patches that also have the removed patch as a dependency. [67] Another problem is inventing a user interface that makes history exploration easy, since traditional commit browsers present exceedingly complex trees.

7.3 Formal Analysis

Precise semantics of current version control systems are defined only loosely and the lack of formal analysis is visible especially in the area of merging algorithms. In order to get a truly reliable and dependable system with regards to merging "wild" branches exhibiting pathological behaviour, the precise semantics of various merges should be specified formally; many of the currently widely used merging methods lack detailed formal analysis and their behaviour in corner cases is only guessed.

There has been ongoing development in this area [20] [2], but it appears to be mostly stalled now as the recursive three-way merge and similar algorithms appear to work well enough in common cases and thus there is not much incentive to formally describe their behaviour in corner cases. However, we believe that such a formal foundation would also help development of newer changeset-oriented systems, in turn making progress with the third-party changes tracking problem.

8 Conclusion

We have gone through the most important concepts, algorithms and structures used in current version control systems, and put them in the context of practical usage; we hope this gave the reader a comprehensive idea about the area and will help to put his further research in concrete topics into the broad context.

Among the open source projects, the new generation of distributed version control systems has registered massive uptake; while there are still projects using CVS and many projects are using or even migrating to SVN, large portion of the most high-profile open source projects with many commits per day and code base of millions of lines of code [43] is using one of the distributed systems described above.

The wide adoption and scale of use clearly indicates that the current systems have reached appropriate scalability and reliability levels. The high number of user interface tools in development also makes it more comfortable to use these systems and smoothens up the learning curve that has been traditionally very steep for distributed version control systems. However, as described in section 7, many desirable features are still unimplemented and support for some workflows is still very lacking.

8.1 Acknowledgements

This article is based on Bachelor thesis I wrote at the Faculty of Math and Physics, Charles University. I would like to thank my thesis advisor Martin Mareš for guidance and Bram Cohen, Piet Delport, Junio C. Hamano, Greg Hudson, Uwe Kleine-König, Tom Lord, Matt Mackall, Larry McVoy, Kenneth C. Schalk, Nathaniel Smith, Pavel Surynek and Zooko Wilcox–O'Hearn for their review and helpful input.

References

[1] Smith N.: *3-way merge considered harmful*, monotone-devel@nongnu.org, 2005-05-01.

[2] Löh A., Swierstra W., Leijen D.: *A Principled Approach to Version Control*, http://people.cs.uu.nl/andres/VersionControl.html

[3] Asklund U., Bendix L.: *A study of configuration management in open source software projects*, Software, IEEE Proceedings **149** (2002) 40–46.

[4] Ziv J., Lepmep A.: *A universal algorithm for data compression*, IEEE Transactions on Information Theory, 23**(3)**:337-343, 1977.

[5] Henson V.: *An Analysis of Compare-by-hash*, USENIX Proceedings of HotOS **IX** (2003).

[6] Stiegler M.: *An Introduction to Petname Systems*, http://www.skyhunter.com/marcs/petnames/IntroPetNames.html

[7] Myers E.: *An O(ND) Difference Algorithm and its Variations*, Algorithmica **1** (1986) 251–266.

[8] Wilcox-O'Hearn Z. B.: *badmerge – concrete version with good semantics*, https://zooko.com/badmerge/concrete-good-semantics.html

[9] Bentley A., Pool M., et al.: *Bazaar*, http://www.bazaar-vcs.org/

[10] Cohen B., et al.: *Bazaar Source Code, bzrlib/patiencediff.py*, http://www.bazaar-vcs.org/Download

[11] McVoy L., et al.: *BitKeeper*, http://www.bitkeeper.com/

[12] Korn M., et al.: *Bzr Virtual Filesystem*, https://launchpad.net/bzr-fs/

[13] Bentley A., Pool M., et al.: *Bzr Weave Format*, http://bazaar-vcs.org/BzrWeaveFormat

[14] Black J.: *Compare-by-Hash: A Reasoned Analysis*, USENIX Annual Technical Conference–Systems and Experience Track, USENIX 2006, Boston, 2006.

[15] Dart S.: *Concepts in configuration management systems*, Proceedings of the 3rd international workshop on Software configuration management, Trondheim, 1991, ISBN 0-897914-429-5.

[16] Torvalds L.: *Concerning Git's Packing Heuristics*, http://www.kernel.org/pub/software/scm/git/docs/technical/pack-heuristics.txt

[17] Grune D., Berliner B., et al.: *CVS*, http://www.nongnu.org/cvs/

[18] Mansfield D., et al.: *CVSps*, http://www.cobite.com/cvsps/

[19] Roundy D.: *Darcs*, http://darcs.net/

[20] Roundy D.: *Darcs manual*, http://darcs.net/manual/

[21] Darcs Community: *Darcs Wiki: Conflicts FAQ*, http://wiki.darcs.net/DarcsWiki/ConflictsFAQ

[22] Tridgell A.: *Efficient Algorithms for Sorting and Synchronization*, PhD thesis, Australian National University, 1999.

[23] Rabin M. O.: *Fingerprinting By Random Polynomials*, Technical Report TR-15-81, Center for Research in Computer Technology, Harvard University, 1981.

[24] Torvalds L., Hamano J. C., et al.: *Git*, http://git.or.cz/

[25] Hamano J. C., et al.: *Git — A Stupid Content Tracker*, OLS 2006 Proceedings Volume 1, Ottawa, 2006.

[26] Greaves D., Hamano J. C., Torvalds L., et al.: *Git Documentation*, http://www.kernel.org/pub/software/scm/git/docs/

[27] Kuivinen F., et al.: *Git Source Code, builtin-merge-recursive.c*, http://www.kernel.org/pub/software/scm/git/

[28] Pitre N., et al.: *Git Source Code, diff-delta.c*, http://www.kernel.org/pub/software/scm/git/

[29] Hamano J. C., et al.: *Git Source Code, diffcore-delta.c*, http://www.kernel.org/pub/software/scm/git/

[30] Torvalds L., et al.: *Git Source Code, pack-objects.c*, http://www.kernel.org/pub/software/scm/git/

[31] Blank M., et al.: *gitfs*, http://www.sfgoth.com/~mitch/linux/gitfs/

[32] Lord T., et al.: *GNU Arch*, http://www.gnu.org/software/gnu-arch/

[33] Eggert P., et al.: *GNU Diffutils*, http://www.gnu.org/software/diffutils/

[34] Henson V., Henderson R., et al.: *Guidelines for Using Compare-by-Hash*, http://infohost.nmt.edu/~val/review/hash2.html

[35] Brownawell T., Smith N.: *Improvements to *-merge*, monotone-devel@nongnu.org, 2005-08-30.

[36] Aldous D., Diaconis P., et al.: *Longest Increasing Subsequences: From Patience Sorting to the Baik-Dieft-Johansson Theorem*, Bull. Amer. Math. Soc. **36** (1999) 413—432.

[37] Mackall M., et al.: *Mercurial*, http://www.selenic.com/mercurial/

[38] Mackall M., et al.: *Mercurial Source Code, mercurial/revlog.py*, http://www.selenic.com/mercurial/release/

[39] Hoare G., Smith N., et al.: *Monotone*, http://monotone.ca/

[40] Hoare G., et al.: *Monotone Documentation: Hash Integrity*, http://monotone.ca/docs/Hash-Integrity.html

[41] Hoare G., et al.: *Monotone FAQ*, http://www.venge.net/mtn-wiki/FAQ

[42] Andrieu O., et al.: *Monotone-viz*, http://oandrieu.nerim.net/monotone-viz/

[43] Wheeler D.: *More Than a Gigabuck: Estimating GNU/Linux's Size*, http://www.dwheeler.com/sloc

[44] Wilcox–O'Hearn Z. B.: *Names: Decentralized, Secure, Human-Meaningful: Choose Two*, https://zooko.com/distnames.html

[45] Hudson G.: *Notes on keeping version histories of files*, http://web.mit.edu/ghudson/thoughts/file-versioning

[46] Shapiro J. S., et al.: *OpenCM*, http://www.opencm.org/

[47] Collins R., et al.: *Patch Queue Manager*, http://bazaar-vcs.org/PatchQueueManager

[48] Ratcliff J. W., Metzener D. E., et al.: *Pattern Matching: The Gestalt Approach*, Dr. Dobb's Journal, July 1988, p. 46.

[49] Seiwald C., et al.: *Perforce*, http://www.perforce.com/

[50] Hudson G., Personal conversation on Subversion

[51] Schalk K. C., Personal conversation on Vesta

[52] IEEE: *POSIX 1003.1 2004 Edition*, 2004.

[53] Python Community: *Python difflib — helpers for producing deltas*, http://docs.python.org/lib/module-difflib.html

[54] Tichy W.: *RCS: a system for version control*, Software—Practice and Experience, **15** (7): 637–654, July 1985.

[55] Research Community: *RevCtrl Wiki*, http://revctrl.org/

[56] Smith N., Baudis P., Cohen B. et al.: *RevCtrl Wiki: Criss Cross Merge*, http://revctrl.org/CrissCrossMerge

[57] National Institute of Standards and Technology: *Secure Hash Standard (SHS)*, FIPS180-02, August 2002.

[58] E-Soft Inc.: *Security Space Survey of Public Subversion DAV Servers*, http://subversion.tigris.org/svn-dav-securityspace-survey.html

[59] Blandy J., Fogel K., et al.: *Subversion*, http://subversion.tigris.org/

[60] Berlin D.: *SVNDIFF version 1*, dev@subversion.tigris.org, 2005-10-21

[61] Cohen B.: *The diff problem has been solved*, http://bramcohen.livejournal.com/37690.html

[62] Cohen B.: *The new Codeville merge algorithm*, revctrl@lists.zooko.com, 2005-05-05.

[63] Rochkind M. J.: *The Source Code Control System*, IEEE Transactions on Software Engineering SE-1:4 (Dec. 1975), 364–370.

[64] Korn D., et al.: *The VCDIFF Generic Differencing and Compression Data Format*, RFC3284.

[65] Heydon A., Levin R., Mann T., Yu Y.: *The Vesta Approach to Software Configuration Management*, Compaq SRC Research Report 168, 2001.

[66] Baudis P., et al.: *TopGit*, http://repo.or.cz/w/topgit.git

[67] Baudis P.: *TopGit: how to deal with upstream inclusion*, git@vger.kernel.org, 2008-09-14.

[68] Mackall M.: *Towards a Better SCM: Revlog and Mercurial*, OLS 2006 Proceedings Volume 2, Ottawa, 2006.

[69] Davison W.: *Unified context diff tools*, comp.source.misc Volume 14, 1990.

[70] Conradi R., Westfechtel B.: *Version models for software configuration management*, ACM Computing Surveys **30** (1998) 232–282.

[71] McCoy K.: *VMS File System Internals*, Digital Press, 1990, ISBN 1-55558-056-4.

[72] MacDonald J.: *XDelta*, http://xdelta.org/

[73] MacDonald J.: *Versioned File Archiving, Compression, and Distribution*, UC Berkeley, 1999.

[74] Trendafilov D., Memon N., Suel T.: *zdelta: An Efficient Delta Compression Tool*, Technical Report TR-CIS-2002-02, Polytechnic University Brooklyn, 2002.